

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 1 122 924 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
08.08.2001 Bulletin 2001/32

(51) Int Cl.7: H04L 29/06

(21) Application number: 00203934.5

(22) Date of filing: 08.11.2000

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: Nazari, Siamak
Arcadia, CA 91006 (US)

(74) Representative: Hanna, Peter William Derek
Peter Hanna Associates
11 Mespil Road
Dublin 4 (IE)

(30) Priority: 02.12.1999 US 454246

(71) Applicant: SUN MICROSYSTEMS, INC.
Palo Alto, California 94303 (US)

(54) Method and apparatus for providing local path I/O in a distributed file system

(57) One embodiment of the present invention provides a system that reduces network traffic during remote file system operations by using a local path between a client and a storage device to bypass transfers through a server. The system operates by receiving a request at the server to access a file located on the storage device: this request is received across a network from an application on a client. Next, the system determines if a local path exists between the client and the storage device containing the file. Note that a local path is independent from a first path between the server and the storage device. If the local path exists, the system sends a bypass request to the client. This bypass request commands the client to satisfy the request by transferring data between the client and the storage device through the local path. In a variation on the above embodiment, the system additionally receives a completion indicator from the client. This completion indicator informs the server that the client successfully completed the request across the local path. In response to the completion indicator, the system updates data structures on the server to indicate that the request has been completed, and then forward the completion indicator to the application on the client. In a variation on the above embodiment, if the local path does not exist, the system completes the request along a normal path by transferring the data between the client and the server through the network, and between the server and the storage device through the first path. After doing so, the system updates data structures on the server to indicate that the request has been completed and sends a completion indicator to the application on the client.

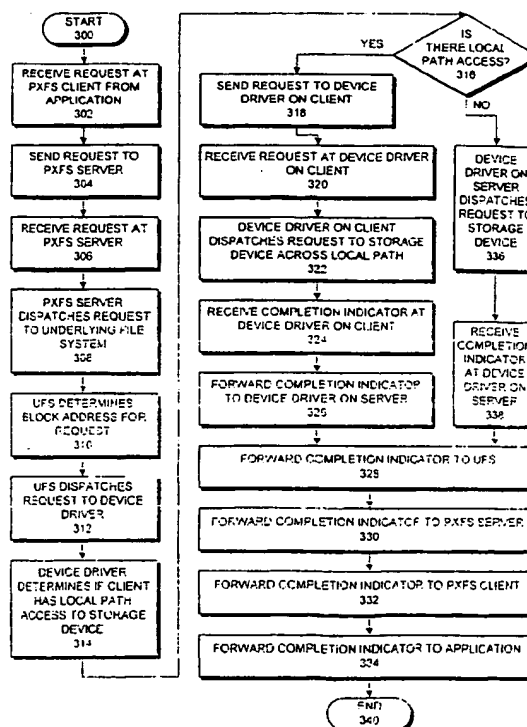


FIG. 3

BEST AVAILABLE COPY

EP 1 122 924 A2

Description**BACKGROUND****Field of the Invention**

[0001] The present invention relates to file systems that are accessible across computer networks. More particularly, the present invention relates to a method and an apparatus for reducing network traffic during remote file system operations by using a local path between a client and a storage device to bypass transfers through a server.

Related Art

[0002] As computer networks are increasingly used to link computer systems together, distributed operating systems have been developed to control interactions between computer systems across a computer network. Some distributed operating systems allow client computer systems to access resources on server computer systems. For example, a client computer system may be able to access a file on a server computer system across a network. Such distributed file systems make it easy to manipulate files located on a remote server. However, if such distributed file systems are not designed carefully, they can easily generate unnecessary communications across the network, which can degrade overall system performance.

[0003] Some distributed operating system have been designed to facilitate high availability. A highly available system provides redundancy so that when a primary node in the highly available system fails a backup secondary node is able to take its place. This allows the system to continue functioning without interruption. Highly available computer systems often provide multiple servers to access a storage device. Each of these servers has its own local data path to the storage device so that if one server fails another can take its place.

[0004] At first glance, there appears to be an advantage in using the multiple paths to the storage device to provide parallel access to the storage device, in order to increase overall computer system performance. However, allowing parallel accesses to the storage device can give rise to concurrency and locking problems which can greatly complicate the design of a distributed operating system. For these reasons, it is easier to centralize control over file system accesses within a single server.

[0005] However, routing communications through a single server can impede system performance. A major bottleneck is caused by communications between the server and a client across a network. Each block of data retrieved from a storage device must be divided into numerous smaller packets, which are transferred over a network between the client and the server. For example, a 64K block of data retrieved from a disk drive is typically divided into dozens of packets, which are individually

sent across the computer network between the client and the server. This process can greatly degrade system performance because each of the packets can generate an interrupt on the server and an interrupt on the client. These interrupts consume a great deal of computational time because process state is saved and restored during each interrupt.

[0006] Note that such data transfers across the computer network are largely wasted in the case where the client computer system possesses its own local path to the storage device. What is needed is a method and an apparatus that reduces network traffic during remote file system operations by using a local path between a client and a storage device whenever possible to bypass transfers through a server.

SUMMARY

[0007] One embodiment of the present invention provides a system that reduces network traffic during remote file system operations by using a local path between a client and a storage device to bypass transfers through a server. The system operates by receiving a request at the server to access a file located on the storage device; this request is received across a network from an application on a client. Next, the system determines if a local path exists between the client and the storage device containing the file. Note that a local path is independent from a first path between the server and the storage device. If the local path exists, the system sends a bypass request to the client. This bypass request commands the client to satisfy the request by transferring data between the client and the storage device through the local path.

[0008] In a variation on the above embodiment, the system additionally receives a completion indicator from the client. This completion indicator informs the server that the client successfully completed the request across the local path. In response to the completion indicator, the system updates data structures on the server to indicate that the request has been completed, and then forwards the completion indicator to the application on the client.

[0009] In a variation on the above embodiment, if the local path does not exist, the system completes the request along a normal path by transferring the data between the client and the server through the network, and between the server and the storage device through the first path. After doing so, the system updates data structures on the server to indicate that the request has been completed and sends a completion indicator to the application on the client.

[0010] In one embodiment of the present invention, the request can be a read request to read the data from the file or a write request to write the data to the file.

[0011] In one embodiment of the present invention, the bypass request to the client includes, a pointer to a buffer on the client, an identifier for the storage device.

a block offset into the storage device, and a length of the request. In this embodiment, if the request is a read operation, the buffer on the client holds the data that is read from the storage device. If the request is a write operation, the buffer on the client holds the data to be written to the storage device.

[0012] One embodiment of the present invention provides a system that reduces network traffic during remote file system operations by using a local Path between a client and a storage device to bypass transfers through a server. The system operates by receiving a request, at the client, to access a file located on the storage device. This request is received from an application running on the client. Next, the system sends the request to the server. The system waits for a request completion indicator to be received from the server. If a bypass request is received from the server, the system transfers data between the client and the storage device through a local path that is independent from a first path between the server and the storage device. Next, the system sends a local path completion indicator to the server. This local path completion indicator informs the server that the data has been transferred across the local path. If a request completion indicator is received from the server, the system forwards the request completion indicator to the application. This request completion indicator specifies that the request has been completed by transferring the data across either the first path or the local path.

BRIEF DESCRIPTION OF THE FIGURES

[0013] FIG. 1 illustrates a distributed computer system in accordance with an embodiment of the present invention.

[0014] FIG. 2 illustrates a layered architecture for a distributed file system in accordance with an embodiment of the present invention.

[0015] FIG. 3 is a flow chart illustrating the process of using a local path between a client and a storage device to bypass transfers through a server in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION

[0016] The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

[0017] The data structures and code described in this

detailed description are typically stored on a computer readable storage medium, which may be any device or medium that can store code and/or data for use by a computer system. This includes, but is not limited to, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact discs) and DVDs (digital video discs), and computer instruction signals embodied in a transmission medium (with or without a carrier wave upon which the signals are modulated). For example, the transmission medium may include a communications network, such as the Internet.

Distributed Computer System

[0018] FIG. 1 illustrates a distributed computer system 100 in accordance with an embodiment of the present invention. Distributed computer system 100 includes client 102 coupled to server 122. Client 102 may include any node on a network including computational capability and including a mechanism for communicating across the network. Server 122 may include any node on a network including computational capability and data storage capability, as well as a mechanism for servicing requests from client 102 for computational and data storage resources.

[0019] Hub 150 implements a network that couples together client 102 and server 122. Client 102 and server 122 use hub 150 to communicate with each other. Hub 150 is a special case of a general network that can be used with the present invention. This general network can include any type of wire or wireless link between computers including, but not limited to, a local area network, a wide area network, or a combination of networks.

[0020] Note that client 102 is coupled to hub 150 through network interface controller (NIC) 106. Similarly, server 122 is coupled to hub 150 through network interface controller 126. Network interface controllers 106 and 126 can include any type of interface for coupling client 102 and server 122 with hub 150.

[0021] FIG. 1 also includes storage device 130, which is coupled to client 102 and server 122. Storage device 130 includes non-volatile storage for code and/or data and operates under control of client 102 and/or server 122. In one embodiment of the present invention, storage device 130 includes a disk drive.

[0022] Storage device 130 is coupled to disk controller 128 within server 122 through path 129. Storage device 130 is also coupled to disk controller 108 within client 102 through local path 109. Disk controllers 108 and 128 include circuitry and software to control a number of disk spindles. Note that local path 109 is independent from path 129.

[0023] Client 102 includes central processing unit (CPU) 104. Similarly, server 122 includes CPU 124. Central processing units 104 and 124 can include any type of processing engines that can be used in a computer system, including, but not limited to, microproces-

sors, mainframe processors, device controllers and processing circuitry within appliances.

[0024] CPU 104 is coupled to memory 105 which includes code to implement local path I/O in accordance with an embodiment of the present invention. CPU 124 is coupled to memory 125, which similarly includes code to implement local Path I/O. Memories 105 and 125 include any type of random access memory containing code and/or data for CPUs 104 and 124, respectively.

[0025] In the embodiment of the present invention illustrated in FIG. 1, client 102 can act as a backup server for fault-tolerance purposes. For example, assume server 122 is the primary server and client 102 functions as a secondary server for storage device 130. If primary server 122 fails during a file system operation, client 102 takes its place. Note, however, that the present invention can be applied to any type of distributed file system, and is not limited to fault-tolerant file systems.

[0026] In one embodiment of the present invention, the system illustrated in FIG. 1 operates under control of the SOLARIS™ MC operating system, which is a product of Sun Microsystems, Inc. of Palo Alto, California. The Solaris MC operating system is a UNIX-based operating system. Hence, in describing the present technology, UNIX terminology and concepts are frequently used. However, this usage is for purposes of illustration and is not to be construed as limiting the invention to this particular operating system.

[0027] Sun, the Sun logo, Sun Microsystems and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Layered Architecture

[0028] FIG. 2 illustrates a layered architecture for a distributed file system in accordance with an embodiment of the present invention. The left-hand side of FIG. 2 illustrates layers that are present on a client side of the distributed file system. The right-hand side of FIG. 2 illustrates layers that are present on a server side of the distributed file system.

[0029] The client side of the distributed file system includes an application 203, which makes file system requests. Application 203 operates in user space 202 and makes file system calls to proxy file system client 205 within kernel space 204. Proxy file system client 205 is the client-side portion of a distributed file system. Proxy file system client 205 communicates with proxy file system server 206 within server 122. Note that this communication takes place across a network through hub 150.

[0030] Proxy file system server 206 is the server-side portion of the distributed file system. Proxy file system server 206 communicates with underlying file system 208 on server 122.

[0031] Underlying file system 208 communicates with device driver 210, which contains code for communicat-

ing with disk controller 128 within server 122. Finally, disk controller 128 controls the operation of storage device 130 through path 129. In one embodiment of the present invention, storage device 130 includes a disk drive.

[0032] In the embodiment of the present invention illustrated in FIG. 1, client 102 includes a parallel local path 109 to storage device 130. Proxy file system client 205 can also act as a proxy file system server. To this end, proxy file system client 205 communicates with underlying file system 207 on client 102. Underlying file system 207 communicates with device driver 209, which contains code for communicating with disk controller 108. Finally, disk controller 128 communicates with storage device 130 through local path 109.

[0033] Note that the process of bypassing file system operations to local path 109 is accomplished through communications between device driver 210 on server 122 and device driver 209 on client 102. This bypassing process is transparent to underlying file system 208 on server 122. Also note that the bypassing process can alternatively be implemented between proxy file system server 206 and proxy file system client 205, or between underlying file system 207 and underlying file system 208.

Process of Bypassing Using Local Path

[0034] FIG. 3 is a flow chart illustrating the process of using local path 109 between client 102 and storage device 130 to bypass transfers through server 122 in accordance with an embodiment of the present invention. The system starts by receiving a request to perform a file system operation at proxy file system client 205 within client 102 (step 302). This request originates from application 203, which is executing on client 102. Proxy file system client 205 dispatches the request to proxy file system server 206 (step 304). Proxy file system server 206 receives the request (step 306) and dispatches the request to underlying file system 208 (step 308). Underlying file system 208 determines the block address of the request (step 310) and dispatches the request to device driver 210 (step 312).

[0035] Device driver 210 determines which client made the request, and then determines if the client has local path access to storage device 130 (step 314). Note that a client has local path access to storage device 130 if it can access storage device 130 without using path 129 between storage device 130 and server 122.

[0036] If client 102 has local path access to storage device 130, device driver 210 on server 122 sends the request to device driver 209 on client 102 (step 318). This request includes a pointer to a buffer on client 102 for storing file operation data. During a write operation, this buffer stores data to be written to storage device 130. During a read operation, this buffer stores data that is retrieved during the file operation. The request also includes an identifier for storage device 130 as well as

a block offset into storage device 130 and a length for the I/O operation.

[0037] Device driver 209 on client 102 receives the request (step 320) and then dispatches the request to storage device 130 through disk controller 108 and across local path 109 (step 322). When storage device 130 completes the request, it sends a completion indicator to device driver 209 on client 102. Device driver 209 receives the completion indicator (step 324) and forwards the completion indicator to device driver 210 on server 122 (step 326). Device driver 210 forwards the completion indicator to underlying file system 208 (step 328), which forwards the completion indicator to proxy file system server 206 (step 330). Proxy file system server 206 forwards the completion indicator to proxy file system client 205 (step 332). Finally, proxy file system client 205 forwards the completion indicator to application 203 to complete the file access operation (step 334).

[0038] If, in step 316, device driver 210 on server 122 determines that there is no local path that can be used to bypass the file transfer operation, device driver 210 dispatches the request through a normal path to storage device 130 through disk controller 128 (step 336). Upon receiving a completion indicator from storage device 130 (step 338), device driver 210 forwards the completion request to underlying file system 208 (step 328). Underlying file system 208 forwards the completion indicator to proxy file system server 206 (step 330). Proxy file system server 206 forwards the completion indicator to proxy file system client 205 on client 102 (step 332). Finally, proxy file system client 205 forwards the completion indicator to application 203 to complete the file access operation (step 334).

[0039] The foregoing descriptions of embodiments of the invention have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art. Additionally, the above disclosure is not intended to limit the invention. The scope of the invention is defined by the appended claims.

Claims

1. A method for reducing network traffic during remote file system operations, the method operating in a system including a client (102) and a server (122) for a storage device (130), the client (102) having a local path (109) to the storage device (130), the method comprising:

receiving a request (306), at the server (122), to access a file located on the storage device (130), the storage device (130) being coupled to the server (122) through a first path (129); wherein the request is received from the client

(102) across a network having been generated by an application (203) on the client; determining, at the server (122), if the local path (109) exists between the client (102) and the storage device (130) containing the file; wherein the local path is independent from the first path between the server (122) and the storage device (130); and if the local path exists, sending a bypass request to the client, the bypass request commanding the client (102) to satisfy the request by transferring data between the client (102) and the storage device (130) through the local path.

2. The method of claim 1, further comprising:

receiving a completion indicator (324) from the client, the completion indicator informing the server (122) that the client (102) successfully completed the request across the local path; updating data structures on the server (122) to indicate that the request has been completed; and forwarding the completion indicator to the application (203) on the client.

3. The method of claim 1 or claim 2, wherein if the local path does not exist, the method further comprises:

completing the request along a normal path by transferring the data between the client (102) and the server (122) through the network and between the server (122) and the storage device (130) through the first path; updating data structures on the server (122) to indicate that the request has been completed; and sending a completion indicator (206) to the application (203) on the client.

4. A method according to any preceding claim,

wherein the bypass request to the client (102) includes,

a pointer to a buffer on the client,
an identifier for the storage device (130),
a block offset into the storage device (130) for the request,
and
a length of the request;

wherein if the request is a read operation, the buffer on the client (102) holds the data that is read from the storage device (130); and wherein if the request is a write operation, the buffer on the client (102) holds the data to be

written to the storage device (130) .

5. A computer readable storage medium storing instructions that when executed by a computer cause the computer to perform a method, for reducing network traffic during remote file system operations, the method operating in a system including a client (102) and a server (122) for a storage device (130), the client (102) having a local path to the storage device (130), according to any one of claims 1 to 4. 5
6. A method for reducing network traffic during remote file system operations, the method operating in a system including a client (102) and a server (122) for a storage device (130), the client (102) having a local path to the storage device (130), the method comprising: 10
 - receiving a request (302), at the client, to access a file located on the storage device (130) ; wherein the request is received from an application (203) running on the client; 20
 - sending the request (304) to the server; and if a bypass request is received (320) from the server, 25
 - transferring data between the client (102) and the storage device (130) through the local path (109), the local path being separate from a first path (129) between the server (122) and the storage device (130) so that the data bypasses the server (122), 30
 - and
 - sending a local path completion indicator (328) to the server, the local path completion indicator informing the server (122) that the data has been transferred across the local path (109). 35
7. The method of claim 6, further comprising if a request completion indicator is received from the server, forwarding (334) the request completion indicator to the application (203); 40
 - wherein the request completion indicator specifies that the request has been successfully completed by transferring the data across either the first path or the local path. 45
8. The method of claim 6 or claim 7, 50
 - wherein the bypass request to the client (102) includes, 50
 - a pointer to a buffer on the client,
 - an identifier for the storage device (130),
 - a block offset into the storage device (130) for the request, and 55
 - a length of the request;

wherein if the request is a read operation, the

buffer on the client (102) holds the data that is read from the storage device (130) ; and wherein if the request is a write operation, the buffer on the client (102) holds the data to be written to the storage device (130) .

9. A computer readable storage medium storing instructions that when executed by a computer cause the computer to perform a method, for reducing network traffic during remote file system operations, the method operating in a system including a client (102) and a server (122) for a storage device (130), the client (102) having a local path to the storage device (130), according to any one of claims 6 to 8. 10
10. An apparatus that reduces network traffic during remote file system operations by using a local path between a client (102) and a storage device (130), comprising: 15
 - a receiving mechanism within the server (122) that receives a request to access a file located on the storage device (130), the storage device (130) being coupled to the server (122) through a first path; 20
 - wherein the request is received from the client (102) across a network having been generated by an application (203) on the client; a bypassing mechanism within the server (122) for bypassing file access requests to the local path; 25
 - wherein the bypassing mechanism is configured to determine if the local path exists between the client (102) and the storage device (130) containing the file, the local path being separate from the first path between the server (122) and the storage device (130) ; and 30
 - wherein if the local path exists, the bypassing mechanism is configured to send a bypass request to the client, the bypass request commanding the client (102) to satisfy the request by transferring data between the client (102) and the storage device (130) through the local path. 35
11. The apparatus of claim 10, wherein the bypassing mechanism is further configured to: 40
 - receive a completion indicator from the client, the completion indicator informing the server (122) that the client (102) successfully completed the request across the local path; 45
 - update data structures on the server (122) to indicate that the request has been completed; and
 - forward the completion indicator to the application (203) on the client. 50

12. The apparatus of claim 10 or claim 11, wherein if the local path does not exist, the bypassing mechanism is configured to:

complete the request along a normal path by
transferring the data between the client (102)
and the server (122) through the network and
between the server (122) and the storage de-
vice (130) through the first path;
update data structures on the server (122) to
indicate that the request has been completed;
and
send a completion indicator to the application
(203) on the client.

13. The apparatus of any one of claims 10 to 12,
wherein the bypass request to the client (102)
includes,

a pointer to a buffer on the client,
an identifier for the storage device (130),
a block offset into the storage device (130) for
the request,
and
a length of the request;

wherein if the request is a read operation, the
buffer on the client (102) holds the data that is read
from the storage device (130); and

wherein if the request is a write operation, the
buffer on the client (102) holds the data to be written
to the storage device (130).

14. An apparatus that reduces network traffic during re-
mote file system operations by using a local path
between a client (102) and a storage device (130),
comprising:

an application (203) processing mechanism
within the client (102) that receives a request
from an application (203) to access a file locat-
ed on the storage device (130);
a file access processing mechanism that is con-
figured to send the request to the server;
a bypass processing mechanism that is config-
ured to,

receive a bypass request from the server,
transfer data between the client (102) and
the storage device (130) through the local
path, the local path being separate from a
first path between the server (122) and the
storage device (130) so that the data by-
passes the server, and
send a local path completion indicator to
the server, the local path completion indi-
cator informing the server (122) that the re-
quest has been completed across the local

path.

15. The apparatus of claim 14, wherein the file access
processing mechanism is further configured to:

receive a request completion indicator from the
server, the request completion indicator speci-
fying that the request has been successfully
completed by transferring the data across ei-
ther the first path or the local path; and to
forward the request completion indicator to the
application (203).

16. The apparatus of claim 14 or claim 15,
wherein the bypass request to the client (102)
includes,

a pointer to a buffer on the client,
an identifier for the storage device (130),
a block offset into the storage device (130) for
the request, and
a length of the request;

wherein if the request is a read operation, the
buffer on the client (102) holds the data that is read
from the storage device (130); and

wherein if the request is a write operation, the
buffer on the client (102) holds the data to be written
to the storage device (130).




FIG. 1

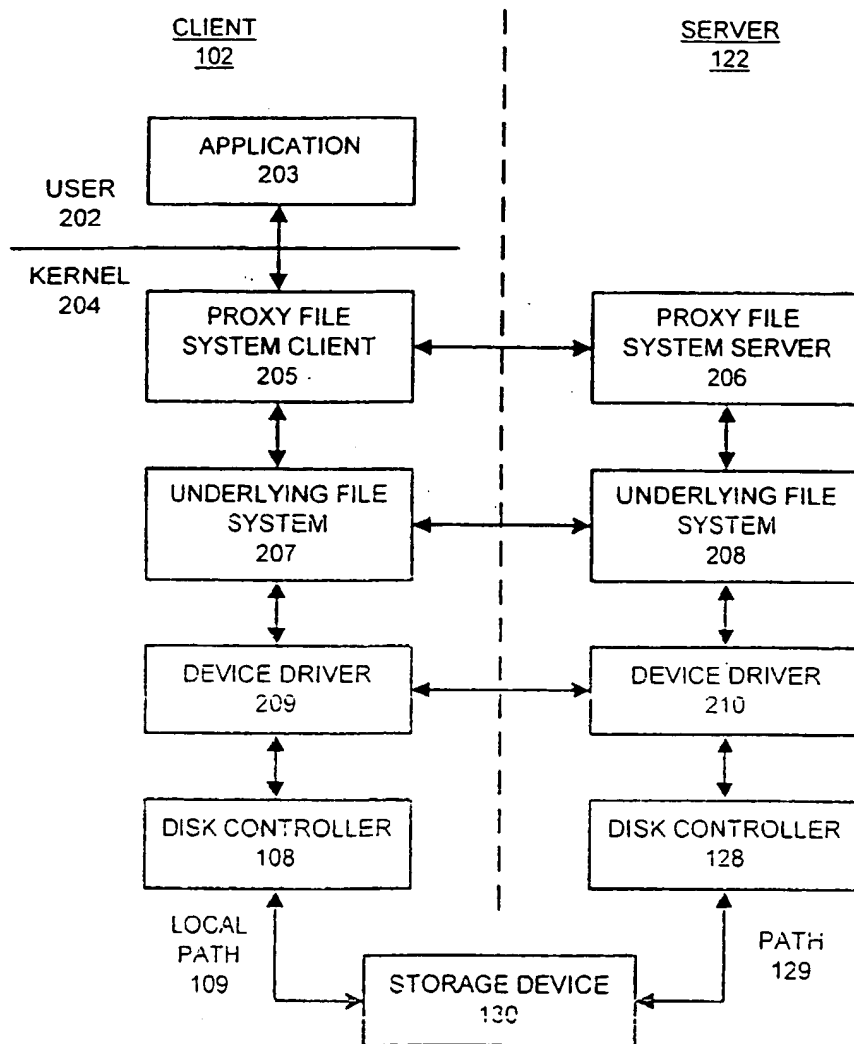


FIG. 2

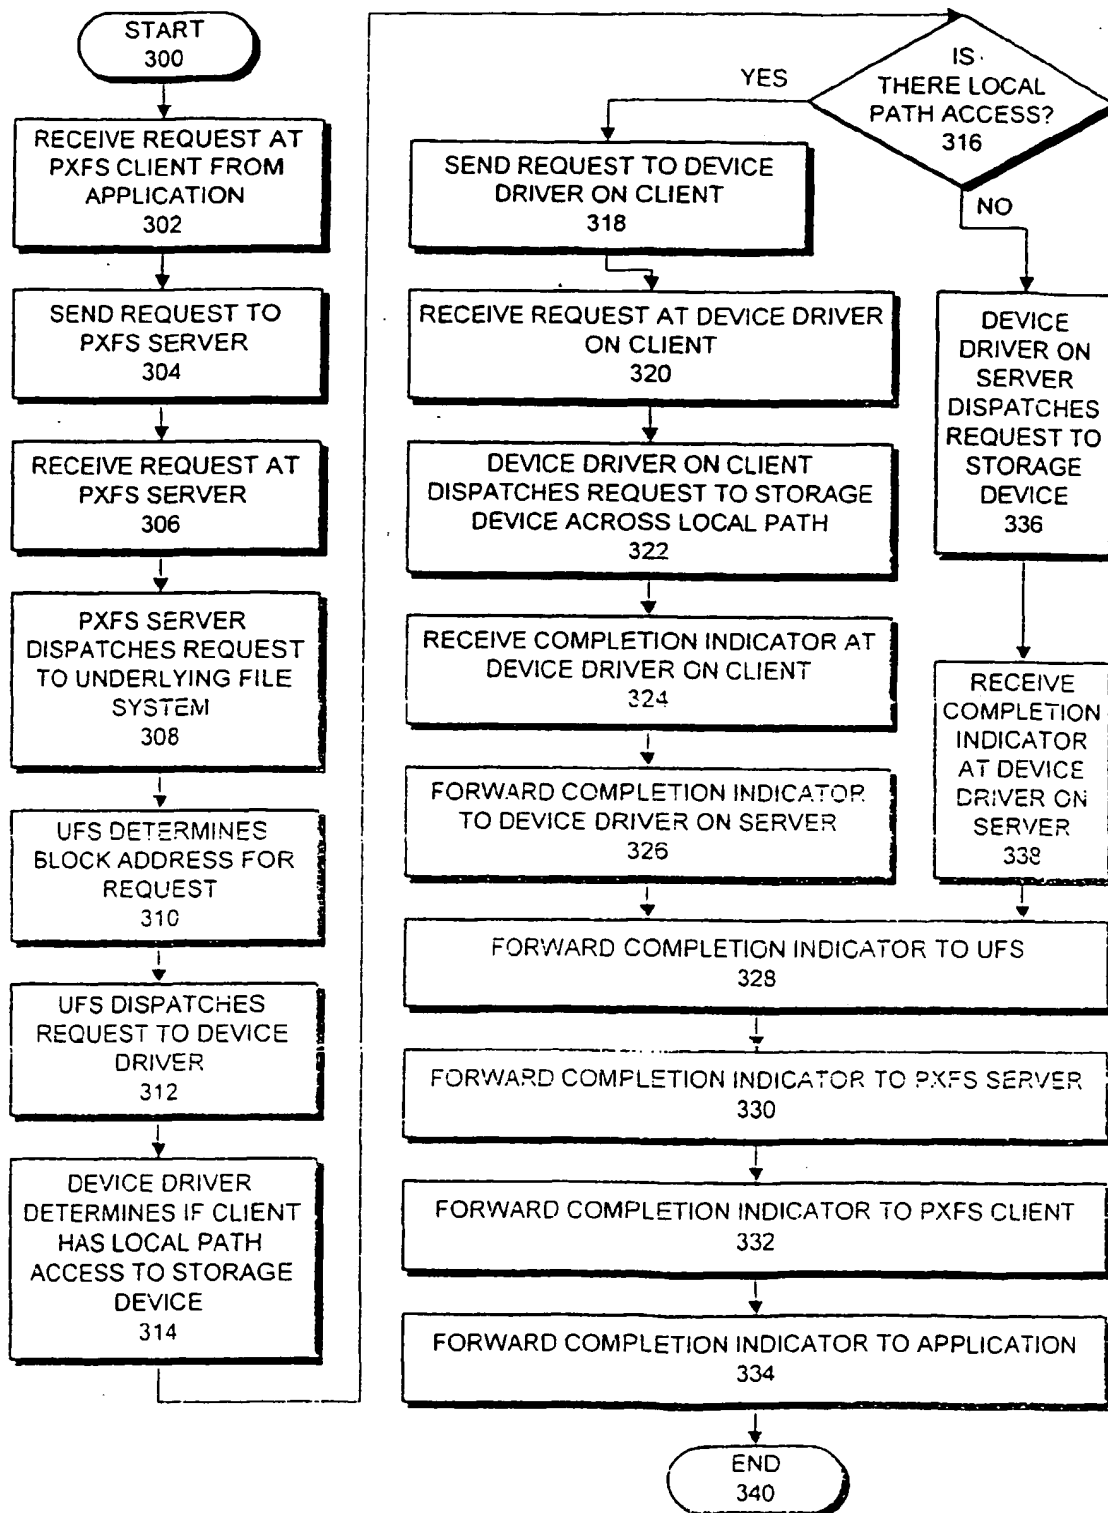



FIG. 3



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 122 924 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
03.03.2004 Bulletin 2004/10

(51) Int Cl.7: H04L 29/06

(43) Date of publication A2:
08.08.2001 Bulletin 2001/32

(21) Application number: 00203934.5

(22) Date of filing: 08.11.2000

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: Nazari, Siamak
Arcadia, CA 91006 (US)

(74) Representative: Hanna, Peter William Derek et al
Hanna, Moore & Curley,
11 Mespil Road,
Dublin 4 (IE)

(30) Priority: 02.12.1999 US 454246

(71) Applicant: SUN MICROSYSTEMS, INC.
Palo Alto, California 94303 (US)

(54) Method and apparatus for providing local path I/O in a distributed file system

(57) One embodiment of the present invention provides a system that reduces network traffic during remote file system operations by using a local path between a client and a storage device to bypass transfers through a server. The system operates by receiving a request at the server to access a file located on the storage device: this request is received across a network from an application on a client. Next, the system determines if a local path exists between the client and the storage device containing the file. Note that a local path is independent from a first path between the server and the storage device. If the local path exists, the system sends a bypass request to the client. This bypass request commands the client to satisfy the request by transferring data between the client and the storage device through the local path. In a variation on the above embodiment, the system additionally receives a completion indicator from the client. This completion indicator informs the server that the client successfully completed the request across the local path. In response to the completion indicator, the system updates data structures on the server to indicate that the request has been completed, and then forward the completion indicator to the application on the client. In a variation on the above embodiment, if the local path does not exist, the system completes the request along a normal path by transferring the data between the client and the server through the network, and between the server and the storage device through the first path. After doing so, the system updates data structures on the server to indicate that the request has been completed and sends a comple-

tion indicator to the application on the client.

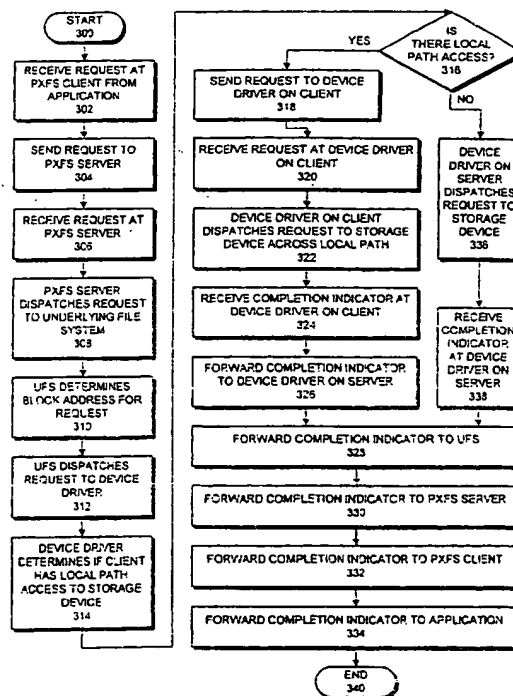


FIG. 3

EP 1 122 924 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 00 20 3934

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	US 5 859 971 A (BITTINGER REED REED ET AL) 12 January 1999 (1999-01-12) * abstract * * column 3, line 10-17 * * column 4, line 1-6 * * column 11, line 3-59 * * column 12, line 1-24 * * figure 2 *	1-16	H04L29/06
X	EP 0 836 145 A (AT & T CORP) 15 April 1998 (1998-04-15) * abstract * * column 2, line 1-20 * * column 3, line 51-58 * * column 4, line 43-56 *	1-16	
A	JP 09 269914 A (FUJITSU LTD) 14 October 1997 (1997-10-14) - & US 2002/059360 A1 (FUJITSU LTD) 16 May 2002 (2002-05-16) * abstract * * paragraphs [0016],[0037] *	2,3,6,7,9,11,12,14,15	
A	US 5 809 527 A (COOPER THOMAS P ET AL) 15 September 1998 (1998-09-15) * abstract * * column 16, line 63 - column 17, line 2 *	4,8,13,16	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 7 January 2004	Examiner Goya, J
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03.02 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 00 20 3934

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

07-01-2004

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5859971	A	12-01-1999	AT 201946 T	15-06-2001
			CA 2218187 A1	21-08-1997
			CN 1184575 A	10-06-1998
			CZ 9703540 A3	18-03-1998
			DE 69613225 D1	12-07-2001
			DE 69613225 T2	08-11-2001
			EP 0823171 A1	11-02-1998
			ES 2159037 T3	16-09-2001
			HU 9801874 A2	30-11-1998
			JP 11500250 T	06-01-1999
			PL 322830 A1	16-02-1998
			WO 9730539 A1	21-08-1997
EP 0836145	A	15-04-1998	US 5931904 A	03-08-1999
			CA 2217370 A1	11-04-1998
			EP 0836145 A2	15-04-1998
			JP 10214239 A	11-08-1998
			US 6240447 B1	29-05-2001
JP 09269914	A	14-10-1997	US 2002059360 A1	16-05-2002
US 5809527	A	15-09-1998	NONE	

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☒ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.